# Bayesian Benchmark Dose Analysis

**Qijun Fang**[a†], **Walter W. Piegorsch**[a,b], and **Katherine Y. Barnes**[a,c]
[a]*Program in Statistics*, [b]*BIO5 Institute*, and [c]*James E. Rogers College of Law*
*University of Arizona, Tucson, AZ, USA*

arXiv:1406.2725v1 [stat.ME] 10 Jun 2014

**Abstract**
An important objective in environmental risk assessment is estimation of minimum exposure levels, called Benchmark Doses (BMDs) that induce a pre-specified Benchmark Response (BMR) in a target population. Established inferential approaches for BMD analysis typically involve one-sided, frequentist confidence limits, leading in practice to what are called Benchmark Dose Lower Limits (BMDLs). Appeal to Bayesian modeling and credible limits for building BMDLs is far less developed, however. Indeed, for the few existing forms of Bayesian BMDs, informative prior information is seldom incorporated. We develop reparameterized quantal-response models that explicitly describe the BMD as a target parameter. Our goal is to obtain an improved estimation and calculation archetype for the BMD and for the BMDL, by employing quantifiable prior belief to represent parameter uncertainty in the statistical model. Implementation is facilitated via a Monte Carlo-based adaptive Metropolis (AM) algorithm to approximate the posterior distribution. An example from environmental carcinogenicity testing illustrates the calculations.

**Keywords**: Adaptive Metropolis sampling, Bayesian modeling, Benchmark analysis, Dose-response analysis, Prior elicitation, Quantitative risk assessment.

# 1 Introduction

## 1.1 Benchmark Risk Analysis

A primary objective in environmental risk assessment is characterization of the severity and likelihood of adverse outcomes caused by a hazardous agent (Stern, 2008). The outcome could be death, cancer, mutation, damage caused by environmental or ecological hazards, etc. In this context, the *risk function*, $R(d)$, is defined as the probability of exhibiting the adverse effect in a subject, object, or system exposed to a particular dose level, $d$, of the hazardous agent. To illustrate, consider the following example.

*Example 1. Benchmarking mammalian carcinogenicity of cumene.*
Cumene, the colloquial name for isopropylbenzene ($C_9H_{12}$), is a hydrocarbon solvent employed in the production of industrial compounds such as phenol and acetone. Occupational and industrial exposures to cumene are common, so the U.S. National Toxicology Program (NTP) explored various forms of mammalian toxicity to the chemical (U.S. NTP, 2009). For example, Table 1 displays quantal-response data on induction of lung tumors (alveolar/bronchiolar adenomas and carcinomas) by cumene in laboratory mice after chronic, two-year, inhalation exposure.

In the table, a clear dose response is evidenced. Of additional interest, however, is characterization of the effects at 'low-dose' exposures, in order to inform risk characterization of this potential carcinogen. We return to these data below.

A contemporary approach to quantal-response risk estimation is known as *benchmark analysis*. First introduced by Crump (1984), the method uses an assumed functional specification for $R(d)$ to

[†]Email address: qijunf@email.arizona.edu; Corresponding author

Table 1: Quantal carcinogenicity data: Alveolar/bronchiolar adenomas and carcinomas in female B6C3F$_1$ mice after two-year inhalation exposure to cumene ($C_9H_{12}$). Source: U.S. NTP (2009).

| Exposure conc. (ppm), $d_i$ | 0 | 125 | 250 | 500 |
|---|---|---|---|---|
| Animals with tumors, $Y_i$ | 4 | 31 | 42 | 46 |
| Animals tested, $N_i$ | 50 | 50 | 50 | 50 |

provide low-dose estimates for the risk. When applied in environmental or public health settings, however, adjustments are made for background or spontaneous effects. Since $R(0)$ represents the background response for all subjects in a population, the risk function is re-expressed in terms of an excess-above-background response rate (Piegorsch and Bailer, 2005, §4.2.1). With quantal data, the *extra risk function* $R_E(d) = \{R(d) - R(0)\}/\{1 - R(0)\}$ is typically employed. From this, the *benchmark dose* (BMD) is calculated by inverting $R_E(d)$ at a predetermined level of risk; the latter is called the *benchmark risk* or *benchmark response* (BMR). In effect, we solve $R_E(\text{BMD}) = \text{BMR} \in (0, 1)$; BMRs between 0.01 and 0.10 are most often seen in practice (U.S. EPA, 2012).

Statistical inferences in benchmark analysis focus on $100(1-\alpha)\%$ confidence (or credible) limits for the BMD. Driven by public health or other safety considerations, only one-sided, lower limits are targeted, denoted as BMDLs (Crump, 1995). Where needed for clarity, we add a subscript for the BMR level at which each quantity is calculated: $\text{BMD}_{100\text{BMR}}$ and $\text{BMDL}_{100\text{BMR}}$. In this fashion, use of BMDs and BMDLs for quantifying and managing risk with a variety of endpoints is growing in both the United States and the European Union (U.S. General Accounting Office, 2001; European Union, 2003; OECD, 2006, 2008).

## 1.2 Parametric Bayesian Benchmark Analysis

With quantal data such as those in Table 1, we take $Y_i$ as the number of responses at the $i^{\text{th}}$ dose level, out of $N_i$ subjects tested at that dose ($i = 1, ..., m$). The standard statistical model here is the binomial, i.e., $Y_i \sim$ indep. $\text{Bin}\big(N_i, R(d_i)\big)$, where $R(d_i)$ is the risk at dose $d_i$. For generic purposes, we denote $\boldsymbol{\theta}$ as an unknown parameter vector that describes $R(d)$.

In risk-analytic studies with quantal data, estimation of the BMD has traditionally been performed via maximum likelihood (Piegorsch and Bailer, 2005, §4.3.1), producing maximum likelihood estimators (MLEs) $\widehat{\boldsymbol{\theta}}$, with corresponding MLEs $\widehat{R}(d)$ for $R(d)$ and $\widehat{R}_E(d)$ for $R_E(d)$. Setting $\widehat{R}_E(d) = \text{BMR}$ and solving for $d$ yields the MLE, $\widehat{\text{BMD}}_{100\text{BMR}}$. The corresponding, frequentist BMDL is built from the statistical properties of $\widehat{\text{BMD}}$.

Appeal to Bayesian modeling and *credible* limits for building BMDLs is far less developed. For quantal-response data, Naufal *et al.* (2009) studied hierarchical dose-response models under a suite of forms available from the U.S. Environmental Protection Agency's Benchmark Dose Software, BMDS (Davis *et al.*, 2012), including the well-known logistic and probit models: $R(d) = \{1 + \exp(-\beta_0 - \beta_1 d)\}^{-1}$ and $R(d) = \Phi(\beta_0 + \beta_1 d)$, respectively. [Here, $\Phi(z)$ is the standard normal cumulative distribution function.] They placed non-informative priors on the dose-response parameters of each model. Shao and Small (2011) also studied the logistic model, along with the quantal-linear model

$$R(d) = 1 - \exp(-\beta_0 - \beta_1 d) \tag{1.1}$$

(where $\beta_0 \geq 0$ and $\beta_1 \geq 0$), a popular dose-response form in risk assessment (Buckley *et al.*, 2009). They similarly placed objective priors on the $\beta$-parameters; further developments appeared in Shao and Small (2012). Shao (2012) expanded these considerations to the probit model, and also introduced a power prior to build historical control information into the hierarchy. Wheeler and Bailer (2012) presented semi-parametric Bayesian models for the dose response, incorporating a probit kernel and cubic B-splines. They placed normal priors on the basis-function coefficients and built dose-response monotonicity into their prior constraints. The normal priors were often flattened to be essentially non-informative (or *objective*) although like Shao (2012), Wheeler and Bailer built informative priors at $d = 0$ to incorporate potential historical control information. Guha *et al.* (2013) described a nonparametric Bayesian model for the quantal setting, with beta/Dirichlet priors on pertinent probabilities related to their nonparametric construction.

The Bayesian approach for BMD estimation has also been applied with *non*-quantal data, and as in the quantal setting its use enjoys increasing application. See, e.g., Wheeler and Bailer (2009b), Morales *et al.* (2006), Held (2004), or Coull *et al.* (2003).

Notably, previous methods for calculating parametric Bayesian BMDs typically assume objective (and sometimes improper) prior distributions for the unknown model parameters. For example, in the logistic model mentioned above the $\beta$-parameters are interpreted as the usual regression intercept and slope on a logit scale. With quantal data, at least, it is unusual for truly informative, risk-analytic, prior knowledge to be available on parameters such as $\beta_0$ and $\beta_1$, since their interpretation is so generic. Such a strategy neglects informative prior information that may be available on the true quantity of interest in this setting, the BMD.

Of course, informative prior knowledge on the BMD may not always exist in practice, and in this case objective priors serve a useful purpose. When such prior information is available, however, the Bayesian paradigm can achieve its true potential. Herein, we explore this by reparameterizing the quantal-response model to explicitly incorporate the BMD. Following on previous presentations for parametric Bayesian benchmark modeling (Shao and Small, 2011, 2012), we focus on the popular quantal-linear dose-response model in (1.1), due to its heavy use in environmental toxicology. Section 2 describes our reparameterization strategy in greater detail. Section 3 develops a Bayesian framework using the reparameterized quantal-linear form for $R(d)$. Section 4 returns to the cumene carcinogenicity data from Example 1, while Section 5 ends with a brief discussion.

# 2 The Quantal-Linear Dose-Response Model

## 2.1 Benchmark Analysis with the Quantal-Linear Model

The quantal-linear form $R(d) = 1 - e^{-\beta_0 - \beta_1 d}$ from (1.1) is a popular construct for dose-response modeling. Also known as the one-stage model or the complementary-log model, it is a special case of the famous multistage model in carcinogenicity testing (Piegorsch and Bailer, 2005, §4.2.1). This has led it to wide application in toxicological risk assessment with quantal data (Foronda *et al.*, 2007; Buckley *et al.*, 2009; Shao and Small, 2012). Notice that $R(0) = 1 - e^{-\beta_0}$ so that the quantal-linear extra risk is simply $R_E(d) = 1 - e^{-\beta_1 d}$. To find the BMD, we solve $R_E(\text{BMD}) = \text{BMR}$ to obtain

$$\text{BMD}_{100\text{BMR}} = -\frac{\log(1 - \text{BMR})}{\beta_1}. \tag{2.1}$$

Recall that we require $\beta_j \geq 0$ $(j = 0, 1)$ in (1.1). [In the exceptional setting where $\beta_1 = 0$, no dose effect is present. If so, the agent under study would be viewed as innocuous, consistent with the corresponding 'definition' of BMD $= \infty$ under (2.1).]

## 2.2 Reparameterizing the Quantal Linear Model

For benchmark risk analysis, a risk assessor, toxicologist, or other domain expert may often have prior knowledge about the target parameter, the BMD, and possibly also about other application-specific values such as the risk at certain doses; cf. the use of control-response information for $R(0)$ by Shao (2012) and Wheeler and Bailer (2012). Our goal is to reparameterize $R(d)$ in terms of meaningful parameters whose prior distributions are more intuitive to elicit in practice. The reparameterization strategy is, of course, not new, even in benchmark analysis; e.g., Parham and Portier (2005, §14.3.4) re-expressed the quantal-linear model in terms of the BMD to facilitate construction of frequentist BMDLs. Following on their lead, we reformulate the unknown regression parameters in (1.1) in terms of well-understood risk-analytic quantities: the target value, BMD, and the background risk, $R(0)$. [The latter quantity is technically a nuisance parameter here, but it is nonetheless likely to be associated with non-trivial prior information, e.g., from historical control databases (Wheeler and Bailer, 2012).] For simplicity, denote these as $\xi = \text{BMD}$ and $\gamma_0 = R(0)$, respectively. Under (1.1), we know $\xi = -\log(1 - \text{BMR})/\beta_1$ and $\gamma_0 = 1 - e^{-\beta_0}$. [For convenience in the sequel, we generally suppress the BMR subscript on $\xi$, although it is understood that the BMD is dependent on the chosen level of $\text{BMR} \in (0, 1)$.] With these, $R(d)$ becomes:

$$R(d) = 1 - \exp\left\{\log(1 - \gamma_0) + \frac{\log(1 - \text{BMR})}{\xi} d\right\}. \tag{2.2}$$

The corresponding extra risk is simply $R_E(d) = 1 - (1 - \text{BMR})^{d/\xi}$. While the notation may appear more burdensome, explicit incorporation of the parameters $\xi$ and $\gamma_0$ allows us to formulate a clearer, more application-oriented hierarchical model. We explore this further in the next section.

# 3 Bayesian Benchmark Analysis

## 3.1 Prior Specification

Under our reformulation the unknown parameter vector is $\boldsymbol{\theta} = [\xi\ \gamma_0]^T$. To construct a Bayesian hierarchy, we assign a joint p.d.f. to $\boldsymbol{\theta}$: $\pi(\boldsymbol{\theta}) = \pi(\xi, \gamma_0)$. Mimicking previous Bayesian models for benchmark analysis (Shao and Small, 2011; Shao, 2012), we assume the unknown parameters enter into the prior independently, so that $\pi(\xi, \gamma_0) = \pi(\xi)\pi(\gamma_0)$. This assumption is not as arbitrary as it may seem: to assess potential correlation between the two quantities, we accessed data from a study (Nitcheva *et al.*, 2007) of 91 different chemical carcinogens archived in the U.S. EPA's Integrated Risk Information System (IRIS: see `http://www.epa.gov/iris`). Each of the 91 IRIS data sets represented a quantal dose-response to a particular carcinogen, allowing us to compute (frequentist) ML estimates for both $\xi$ and $\gamma_0$ under (1.1). We then calculated the Spearman rank correlation $r_S$ (Kvam and Vidakovic, 2007, §7.3) across the 91 pairs of MLEs $(\widehat{\xi}, \widehat{\gamma_0})$. We found the correlation to be insignificantly different from zero: $r_S = -0.039; P > 0.20$.

To specify the individual components of the joint prior, we employ established, flexible forms. For $\xi$ we make an inverse gamma assumption: $\xi \sim IG(\alpha, \beta)$ with marginal prior density $\pi(\xi|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \xi^{-(\alpha+1)} e^{-\beta/\xi} I_{(0,\infty)}(\xi)$, where $\Gamma(a)$ is the usual gamma function and $I_{\mathbb{A}}(x)$ is the indicator function that returns 1 if $x \in \mathbb{A}$ and 0 otherwise. For $\gamma_0 = R(0)$ we take $\gamma_0 \sim Beta(\psi, \omega)$ with marginal prior $\pi(\gamma_0|\psi, \omega) = \frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)} \gamma_0^{\psi-1} (1 - \gamma_0)^{\omega-1} I_{(0,1)}(\gamma_0)$. (Below, we examine how deviations from these priors—emphasizing sensitivity to the inverse gamma assumption—affect the eventual inferences on $\xi$.) Conceptually, this approach shares similarities with a Bayesian inverse dose

estimation strategy suggested by Babb *et al.* (1998) for use in clinical trials; however, technical differences exist between the two methods, driven primarily by dissimilarities between their medical-safety application and our environmental-risk motivation.

The various hyperparameters, $\alpha$, $\beta$, $\psi$, and $\omega$, require specification for implementation of our model. (One could also build further levels into the hierarchy by constructing hyperprior p.d.f.s, and this has potential for future study. We do not investigate it here, however.) Our goal is elicitation, to the best extent possible, of each marginal prior by incorporating the domain expert's prior knowledge.

A broad literature exists on how to conduct prior elicitation; see, e.g., O'Hagan *et al.* (2006) or Kuhnert (2011) and the references therein. From it, we adopt the general strategy that domain experts with minimal statistical expertise are best able to provide prior information in the form of basic location summaries. Medians and other quartiles can be especially effective: based on interactions with environmental toxicologists and risk assessment domain experts, we have found that for the target parameter $\xi$, specification of the the median ($Q_2$), along with the first/lower quartile ($Q_1$) of the IG prior is most effective. Put simply, prior expert knowledge for $\xi$ is likely to be more accurate closer to the origin, since BMDs are associated with adverse effects at low doses. Similarly, for the beta prior on $\gamma_0$ we also elicit the two quartiles $Q_1$ and $Q_2$. We then solve for the pertinent hyper-parameters given these quartiles. (Details are provided in a Supplemental Document.) Of course, other elicitation values are possible, and analysts may wish to experiment with selection of other quartiles, terciles, percentiles, etc.

Given the resulting values of $\alpha$, $\beta$, $\psi$, and $\omega$, the joint prior for $\boldsymbol{\theta} = [\xi \ \gamma_0]^T$ is fully specified, and we proceed with calculation of the posterior p.d.f. We describe this in the next subsection. Before continuing, however, we acknowledge that cases can arise where the prior elicitation breaks down, say, if prior experience with an environmental agent is so limited that the toxicologist simply has no idea where the BMD will lie. When this occurs, moving to objective specifications for the prior densities may be necessary. Many possibilities exist for building objective priors on a strictly positive quantity such as $\xi$, and on a probability such as $\gamma_0$. We favor a simple approach: for an objective prior on $\xi$, use $\xi \sim IG(0.001, 0.001)$ [another option could be $\xi \sim Gamma(0.001, 0.001)$; see below]. The IG prior is a popular suggestion in the literature for right-skewed, positive quantities (Lambert *et al.*, 2005; Christensen *et al.*, 2011, §1.2), such as the BMD. One could alternatively appeal to an (improper) objective prior of the form $\pi(\xi) \propto 1/\xi$, a conventional choice that is often recommended for positive parameters (O'Hagan, 1994, §9.17). In order to facilitate the posterior computations, however, we prefer to employ proper prior densities. In fact, the $IG(0.001, 0.001)$ prior can approximate the improper reciprocal prior quite well when $\xi$ is not extremely close to zero. The $Gamma(0.001, 0.001)$ can also serve as an approximator for $1/\xi$, although for large dose values the approximation can break down. Indeed, we find that scaling the doses to make the highest dose equal 1 is a convenient device that allows both $IG(0.001, 0.001)$ and $Gamma(0.001, 0.001)$ to approximate the conventional improper $1/\xi$ prior. As such, we perform our calculations with doses scaled so that the maximum administered dose equals 1.

For an objective prior on $\gamma_0$, we consider the univariate Jeffreys prior: $\gamma_0 \sim Beta\left(\frac{1}{2}, \frac{1}{2}\right)$. This assigns symmetrically high density to values of $\gamma_0 \to 0$ and $\gamma_0 \to 1$. The former is reasonable in our toxicological setting, although the latter is problematic, since $\gamma_0$ represents the probability of an adverse response under no exposure and is usually small in risk-analytic applications. (A similar complaint could be raised against use of a 'vague' uniform prior for $\gamma_0$.) In the end, even the valid-but-imprecise argument that low prior weight should be assigned to values of $\gamma_0 \to 1$ is itself a form of subjective prior knowledge, and construction of truly objective priors for $\gamma_0$ may be difficult. As the background response rate in the subject organism under study, $\gamma_0$ will be a

well-understood quantity in practice, and as others have noted there will often be at least some useful historical control information available for it. We expect that calls for objective priors on $\gamma_0$ in this setting will be rare. If in the extreme this is not the case, we default to the objective Jeffreys prior $\gamma_0 \sim Beta\left(\frac{1}{2}, \frac{1}{2}\right)$, although we recognize the enigmatic aspects of such a strategy.

## 3.2 Posterior Analysis via Stochastic Approximation

Given our prior specification from §3.1 under a binomial likelihood, the joint posterior p.d.f. for $\boldsymbol{\theta} = [\xi \ \gamma_0]^T$ is

$$\pi(\xi, \gamma_0|\boldsymbol{Y}) = \frac{\prod_{i=1}^{n} \binom{N_i}{Y_i} R(d_i)^{Y_i}\{1 - R(d_i)\}^{N_i - Y_i}}{m(\boldsymbol{Y})} \frac{\beta^\alpha e^{-\beta/\xi}}{\Gamma(\alpha)\xi^{\alpha+1}} \frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)} \gamma_0^{\psi-1}(1-\gamma_0)^{\omega-1} \tag{3.1}$$

over $\xi > 0$ and $\gamma_0 > 0$. The denominator of $\pi(\xi, \gamma_0|\boldsymbol{Y})$ contains the marginal likelihood

$$m(\boldsymbol{Y}) = \int_0^\infty \int_0^1 \prod_{i=1}^{n} \binom{N_i}{Y_i} R(d_i)^{Y_i}\{1 - R(d_i)\}^{N_i - Y_i} \pi(\xi|\alpha, \beta)\pi(\gamma_0|\psi, \omega) d\gamma_0 d\xi,$$

which is intractable under our elicited prior structure. To evaluate (3.1) we therefore turn to Monte Carlo posterior approximations using an adaptive Metropolis (AM) strategy (Andrieu and Thoms, 2008). This is detailed in the Supplementary Document.

For a given data ensemble $(\boldsymbol{Y}, \boldsymbol{N}, \boldsymbol{d})$, we employ the AM algorithm to produce a bivariate chain of draws, $\{\xi_k, \gamma_{0k}\}_{k=1}^K$. For practical purposes, however, we first check the acceptability of the data before employing the algorithm. Our experience shows that a shallow dose response can create unstable frequentist $\hat{\xi}_{100\text{BMR}}$ estimates, and very shallow responses may cause the model fit to fail entirely. Indeed, the EPA's BMDS software program will not estimate a BMD for a flat or negatively trending dose response (Wheeler and Bailer, 2009a). For these reasons, we chose to mark flat or decreasing-trend data as 'data failures' and not perform estimation for them (see the Appendix).

We apply the AM algorithm to any data set passing our data-failure screen. We also include a 'burn-in' diagnostic—detailed in the Supplementary Document—to ensure that the Monte Carlo chain produces a stable approximation to the posterior distribution. If the diagnostic fails after multiple generations of the AM chain, we consider this an 'algorithm failure' and do not report values for the BMD or BMDL. (This is a rare occurrence. Again, see the Supplementary Document for specifics.) The larger the value of $K$ and the more stable the resulting chain, as determined by the diagnostics mentioned above, the better the approximation. Our experience suggests that a starting value of $K = 100,000$, with reduction for burn-in to about 90,000 or 80,000 draws, generally provides stable results. We derive inferences on $\xi$ from this retained sample of draws.

## 3.3 Bayesian Estimation and Inference on the BMD

For estimating the BMD, we appeal to standard decision-theoretic principles and select the Bayes estimator based on minimization of the Bayes risk, after specification of an underlying loss function. As is well-known, choice of squared-error loss leads to the posterior mean, $\mathrm{E}[\xi|\boldsymbol{Y}]$, as the Bayes estimator, while absolute-error loss leads to the posterior median (Casella and Berger, 2002, §7.3.4). We can approximate these posterior quantities using our Monte Carlo sample's arithmetic mean or median, respectively. If subject-matter considerations cannot guide the choice of loss function, and thus which estimator to employ, appeal to absolute-error loss and the posterior

median might be preferable. This suggestion is, admittedly, empirical: with small sample sizes we generally find the (approximated) posterior for $\xi$ to exhibit a right skew, and the median is more robust against large skews. This has an important, practical consequence: note that a larger estimated BMD implies a higher level of 'acceptable' exposure to a potential toxic agent. If this estimated value is artificially inflated due to anomalies such as a heavy right skew, any public health or environmental guidelines based on the estimate may be unnecessarily lax, and even unsafe. Using the posterior median rather than the posterior mean in such a situation would represent a more-precautionary course of action.

Alert readers may notice that from a decision-theoretic perspective, this issue of how to estimate the BMD begs a larger question: both the squared-error and absolute-error losses treat deviations symmetrically. Should we employ an *asymmetric* loss function here instead? Arguably, yes. In effect, a high BMD quantitatively views large exposures to a potentially hazardous agent as relatively safe. If, in truth, the agent is highly toxic, the consequences of such a decision could be severe from a public health or environmental safety perspective. This is more consequential than incorrectly driving BMD $\to 0$ and imposing strict exposure limitations on an innocuous or weakly toxic agent. Unless (considerable, we would contend) socioeconomic factors can counterbalance these safety concerns, BMD overestimation generates a greater 'loss'.

Many asymmetric loss constructions are possible; we employ the simple bilinear loss function of O'Hagan (1994, §2.46):

$$L(\Delta, \xi) = \begin{cases} a(\xi - \Delta) & \Delta \leq \xi \\ b(\Delta - \xi) & \Delta > \xi \end{cases} \tag{3.2}$$

where $\xi$ is the target quantity, estimated by the decision $\Delta$. The constants $a$ and $b$ tune the bilinear loss for each individual application. (When $a = b$ we recover absolute-error loss.) Since we treat overestimation of $\xi$ more harshly than underestimation, we take $b > a > 0$. O'Hagan (1994) shows that under (3.2), the optimal Bayes estimator for $\xi$ is the $100\left(\frac{a}{a+b}\right)$th percentile of the posterior distribution $\pi(\xi|\boldsymbol{Y})$. As a default choice for benchmark dose estimation, we suggest setting the ratio $\frac{a}{b} = \frac{1}{2}$, i.e., overestimation of $\xi$ incurs twice as much relative loss as underestimation. If so, the optimal Bayesian estimator $\widehat{\xi}_{100\text{BMR}}$ becomes the lower/first tercile of $\pi(\xi|\boldsymbol{Y})$. We estimate this with the lower tercile from our Monte Carlo sample of $\xi$.

For a Bayesian BMDL, say, $\underline{\xi}_{100\text{BMR}}$, we essentially desire a one-sided, lower, $100(1-\alpha)\%$ credible limit on $\xi$, satisfying $P(\xi > \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}) = 1 - \alpha$. At the traditional level of $\alpha = 0.05$, this is the lower $5^{\text{th}}$ percentile of $\pi(\xi|\boldsymbol{Y})$, which we approximate via the lower $5^{\text{th}}$ percentile from our Monte Carlo chain.

## 3.4 Prior Sensitivity: $\epsilon$-Contamination Analysis

To investigate the influence of our IG prior assumption on the eventual BMDL $\underline{\xi}_{100\text{BMR}}$, we consider a prior sensitivity analysis. (One could also perform a sensitivity analysis of the beta prior for $\gamma_0$; however, since $\gamma_0$ is an obvious nuisance parameter in terms of BMD estimation and inference, we do not highlight that alternative here.) Our experiences with $\xi$ for this model show that the parameter uncertainty can almost always be described via a right-skew, which helps motivate the IG prior specification. As mentioned above, however, an obvious alternative is the traditional gamma distribution: $\xi \sim \mathit{Gamma}(\alpha, \beta)$, where $\alpha$ and $\beta$ are elicited using the same expert opinion as detailed above.

We follow O'Hagan (1994, §7.15) and rewrite the prior density on $\xi$ as a contaminated mixture of IG and gamma priors. Treating the IG density as the base prior, $\pi_0(\xi)$, and the gamma as the

contaminating density, $q(\xi)$, the prior for $\xi$ is written as

$$\pi(\xi) = (1-\epsilon)\pi_0(\xi) + \epsilon q(\xi), \tag{3.3}$$

where $\epsilon \in [0,1]$ controls the degree of prior contamination. When $\epsilon = 0$, no contamination appears and the prior for $\xi$ is the base IG density; when $\epsilon = 1$, the prior density function for $\xi$ is completely replaced by the contaminating gamma prior.

Employing the prior density in the $\epsilon$-contamination form from (3.3), we denote $\underline{\xi}_{100\text{BMR}}(\epsilon)$ as the estimated BMDL for a specific $\epsilon$. We then monitor the evolution of $\underline{\xi}_{100\text{BMR}}(\epsilon)$ as $\epsilon$ increases from 0 to 1.

An additional quantification of the $\epsilon$-contaminating prior's influence is the instantaneous rate of change of the posterior inference at $\epsilon = 0$ (O'Hagan, 1994, §7.15). We calculate this as the (absolute value of the) first derivative of $\underline{\xi}_{100\text{BMR}}(\epsilon)$ at $\epsilon = 0$: $|D(q)| = \left|\lim_{\epsilon\to 0} \frac{1}{\epsilon}\{\underline{\xi}_{100\text{BMR}}(\epsilon) - \underline{\xi}_{100\text{BMR}}(0)\}\right|$. O'Hagan showed that $D(q)$ simplifies to

$$D(q) = \left\{\underline{\xi}_{100\text{BMR}}(1) - \underline{\xi}_{100\text{BMR}}(0)\right\} \frac{m_q(\boldsymbol{Y})}{m_0(\boldsymbol{Y})}, \tag{3.4}$$

where $m_q(\boldsymbol{Y})$ and $m_0(\boldsymbol{Y})$ are the marginal likelihoods under the contaminating gamma prior density, $q(\xi)$, and base IG prior density, $\pi_0(\xi)$, respectively.

In order to calculate $|D(q)|$, values of $\underline{\xi}_{100\text{BMR}}(1)$ and $\underline{\xi}_{100\text{BMR}}(0)$ are obtained as the estimated BMDLs using $q(\xi)$ and $\pi_0(\xi)$, respectively. To find the marginal likelihoods, $m_q(\boldsymbol{Y})$ and $m_0(\boldsymbol{Y})$, we appeal to our AM sample and employ a geometric estimator via the bridge sampling method recommended by Meng and Wong (1996) and Lopes and West (2004). (Details are given in the Supplementary Document.) $|D(q)|$ can then be used as a measure of sensitivity: as it draws closer to zero, the posterior inferences are less affected by the contaminating prior (O'Hagan, 1994, §7.15). We apply $|D(q)|$ as part of our larger Bayesian benchmark analysis. The next section illustrates these various calculations.

## 3.5 Posterior Calculations on the Extra Risk Scale

Despite its moniker, the benchmark dose is critically related to the extra risk function $R_E(d)$. As West *et al.* (2012) emphasized, the BMD and in particular the BMDL are used to indicate where exposures to hazardous stimuli lead to extra risks at or below the targeted BMR, since they serve as points of departure for an environmental risk assessment. The construction of a joint posterior for $\xi$ and $\gamma_0$ allows for convenient exploration of features on the extra risk scale, by querying their corresponding draws from the bivariate AM-chain.

### 3.5.1 Posterior extra risks at selected doses

Under our reparameterization, a one-to-one correspondence will typically exist between the pair $(\gamma_0, \xi)$ and the $R_E(d)$ curve in the extra risk space. For example, with the quantal-linear model the extra risk $R_E(d) = 1 - (1 - \text{BMR})^{d/\xi}$ is clearly dependent upon $\xi$ (and is independent of $\gamma_0$, although this is not guaranteed in general): as $\xi$ increases, the quantal-linear $R_E(d)$ curve becomes less concave and drops towards zero in a consistent fashion. Such correspondences allows us to explore the posterior distribution of the extra risk at any given $d$, after obtaining the joint posterior distribution of $(\gamma_0, \xi)$.

For example, a risk assessor may be interested in the posterior distribution of the extra risk at certain crucial or regulatory dose levels. Given a posterior AM sample of $(\gamma_0, \xi)$ pairs, sans burn-in, one can map these to the corresponding posterior sample of $R_E(d)$ curves. Then, the vertical

intersection of this sample of curves at various dose levels $d$ corresponds to a posterior sample of the extra risk at that $d$. Density estimates of this posterior can be constructed via standard kernel methods, allowing the analyst to compare the posterior densities across a variety of pertinent doses; the latter could include the BMDL, the BMD estimate (or estimates, if different selection criteria are under consideration), etc. We explore this possibility with the cumene data in §4.3, below.

### 3.5.2 Simultaneous credible bands

This correspondence between $(\gamma_0, \xi)$ pairs and $R_E(d)$ curves open up a panoply of inferences on the extra risk scale. For instance, a posterior 'point' estimate of $R_E(d)$ corresponds to a point estimate of $(\gamma_0, \xi)$. The simplest example would be calculation of the mean centroid $(\bar{\xi}, \bar{\gamma}_0)$ from the retained AM chain, since it is a well-defined quantity. (Bivariate medians are not always unambiguously defined. And, the sort of asymmetric interest described above for estimating the BMD is not as crucial for estimating the extra risk. Thus a 'central' value such as a mean-based centroid would be a reasonable default.) The curve in $R_E$-space corresponding to $(\bar{\xi}, \bar{\gamma}_0)$ serves as an estimator, $\hat{R}_E(d)$, for the extra risk.

Further, a $1-\alpha$ *simultaneous credible band* can be constructed in a straightforward fashion. Simply identify a central subset of the retained AM chain containing $100(1-\alpha)\%$ of the $(\gamma_0, \xi)$ pairs and sweep out the envelope in $R_E$-space to which these points correspond.

Risk-analytic operations generally focus attention on strictly upper limits for the extra risk—corresponding to lower limits on the BMD. This leads to one-sided, upper, credible bands, and simplifies the band construction. Simply find the 'smallest' retained $(\gamma_0, \xi)$ AM-pair above which $100(1-\alpha)\%$ of the AM-draws lie. The corresponding $R_E(d)$ curve will serve as an upper band on the extra risk, valid simultaneously over all $d \geq 0$. How to define 'smallest' is, of course, open to consideration. As an initial proposal, we suggest recognizing $\gamma_0$'s nuisance status and focusing on how $\xi$ varies throughout $(\gamma_0, \xi)$-space: find the smallest value of $\xi$ above which $100(1-\alpha)\%$ of the $(\gamma_0, \xi)$ AM-pairs lie. [One can view this as a horizontal supporting line—or a supporting hyperplane for models with more than two parameters—separating the upper $100(1-\alpha)\%$ of the retained chain from the lower $100\alpha\%$ in the $(\gamma_0, \xi)$ plane.] In effect, this reduces attention to the marginal posterior for $\xi$ and isolates the $(\gamma_0, \xi)$ pair at $\xi = \underline{\xi}_{100\text{BMR}}$. The corresponding curve in $R_E$-space serves as the $1-\alpha$ upper band.

# 4 Cumene Carcinogenicity Data, Revisited

## 4.1 Benchmark data analysis

We applied our Bayesian approach with the reparameterized quantal linear model to the cumene carcinogenicity data in Example 1. Notice that the $C_9H_{12}$ exposure dose, $d$, is actually a concentration (in ppm) here, and so technically we will compute benchmark concentrations (BMCs) based on the quantal carcinogenicity data. We operated with BMR set to the traditional level of 0.10 (U.S. EPA, 2012). Using input from domain experts, the prior elicitation was based on existing background from the toxicological literature. Table 2 summarizes the elicited values. (Details are given in the Supplemental Document.) The consequent prior distributions employed in the analysis were $\xi \sim IG(0.53, 0.13)$ and $\gamma_0 \sim Beta(1.36, 12.31)$. We then applied our AM approach to approximate the posterior for $\xi$ and $\gamma_0$. No 'data failures' or 'algorithm failures' were encountered.

Table 2: Prior elicitation summary for cumene carcinogenicity data.

| Parameter | Lower quartile, $Q_1$ | Median, $Q_2$ | Prior parameters | |
|---|---|---|---|---|
| $\xi$ | 90 ppm (0.18*) | 250 ppm (0.50*) | $\alpha = 0.53$ | $\beta = 0.13$ |
| $\gamma_0$ | 0.04 | 0.08 | $\psi = 1.86$ | $\omega = 12.31$ |

* after scaling to unit interval

We instituted an AM chain size of $100,000$, from which our convergence diagnostic procedure recommended an initial burn-in of $10,000$ draws. The remaining $90,000$ draws were then employed as the posterior approximation for these data. Trace plots showed adequate mixing of the chain; see the Supplemental Document. Figure 1 displays the corresponding histogram (with overlayed kernel density estimate) for the $\xi$ component. A unimodal and slightly right-skewed marginal posterior distribution for $\xi$ is indicated.

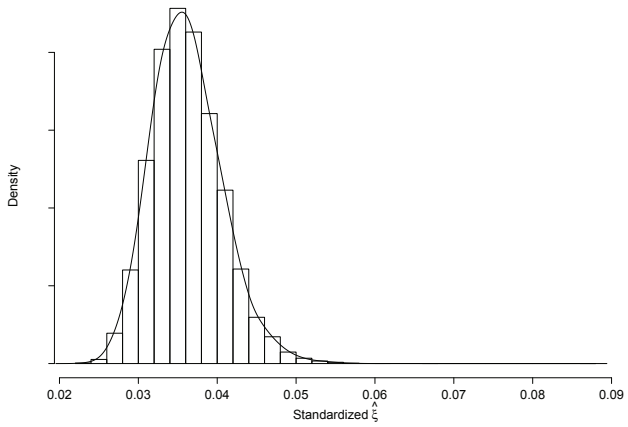


Figure 1: Histogram and kernel density estimator for posterior approximation to $\xi_{10}$ in Example 1. Dose scale ($\xi$) is standardized to unit length after dividing by the highest dose in the data set.

Figure 2 displays the estimated risk function $\hat{R}(d)$ based on the median posterior estimates for $(\xi, \gamma_0)$, the lower tercile estimates for $(\xi, \gamma_0)$, and the MLEs $(\hat{\xi}, \hat{\gamma}_0)$, along with the original observed proportions. All three curves give reasonable estimates for $R(d)$, relative to the data, although the median-based Bayesian estimate shifts consistently to the right of the other two risk functions at most levels of $d$.

For reporting purposes, the final benchmark estimates (when translated back to the original ppm scale) are $\hat{\xi}_{10}$=17.973 ppm if using the sample posterior median or $\hat{\xi}_{10}$=17.046 ppm if using the sample posterior lower tercile. The 95% BMCL is $\underline{\xi}_{10} = 14.752$ ppm. Comparing these to standard frequentist estimates, the MLE is $\hat{\xi}_{10}$=17.062 ppm and a 95% frequentist Wald BMCL is $\underline{\xi}_{10} = 13.618$ ppm. Both sets of values rest in similar ranges, and provide comparable points of departure for conducting further risk-analytic calculations on cumene carcinogenicity.

Thus for these data, our Bayesian approach operates similarly to the frequentist analysis, but

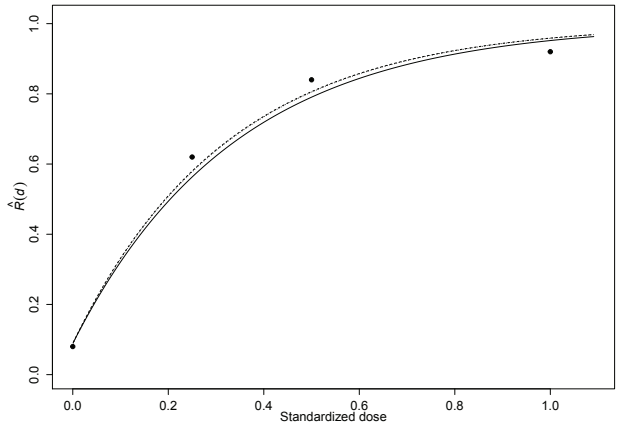


Figure 2: Estimated risk functions for cumene carcinogenicity data in Example 1. Solid curve (———) is based on the posterior median, dashed curve (- - - - -) is based on the lower posterior tercile, and dotted curve (· · · · · ·) is from the maximum likelihood estimate (MLE). The estimated risk function based on the lower tercile is indistinguishable from the estimated ML risk function at this scale. Solid circles are observed proportions. Dose scale ($\xi$) is standardized to unit length after dividing by the highest dose in the data set.

it also provides an additional benefit for risk assessors: the Bayesian strategy combines elicited prior information on both the background risk and the BMC, potentially improving the estimation process by incorporating more-complete prior information.

## 4.2 Prior sensitivity

To explore sensitivity of the IG prior for these data, we returned to the methods in §3.4 and contaminated the base prior for $\xi$ via three scenarios. In Scenario 1, we chose an objective $IG(0.001, 0.001)$ as the base prior for $\xi$ and contaminated it by the similar, objective, $Gamma(0.001, 0.001)$ prior. This scenario was used to investigate the sensitivity of the $IG(0.001, 0.001)$ prior for $\underline{\xi}_{100\text{BMR}}$ when no elicitation is available. In Scenario 2, we chose the elicited IG prior as the base prior for $\xi$ and contaminated it by a similarly elicited gamma prior. We used the same quartile information for $\xi$ to build both elicited priors. Scenario 2 was used to investigate the sensitivity of the IG prior when elicitation is available. In Scenario 3, we chose the elicited IG prior as the base prior for $\xi$ and contaminated it by an objective $Gamma(0.001, 0.001)$ prior. Scenario 3 was used to investigate the robustness of $\underline{\xi}_{100\text{BMR}}$ when the elicited IG prior is contaminated by an objective gamma prior. Throughout, the prior for $\gamma_0$ was taken as either the elicited beta prior, $Beta(1.36, 12.31)$, or an objective $Beta\left(\frac{1}{2}, \frac{1}{2}\right)$ prior. This gave six different settings for study, under which we monitored the consequent BMDL $\underline{\xi}_{100\text{BMR}}(\epsilon)$ as $\epsilon$ increased from 0 to 1. The corresponding $|D(q)|$ instantaneous change measures from (3.4) were also calculated.

Figure 3 displays the evolution of $\underline{\xi}_{100\text{BMR}}(\epsilon)$ across these six settings. The three choices for $\xi$ are distinguished using different line types (see the figure legend), and the two beta priors are distinguished using grey or black shading. In Scenario 1, no substantive change in $\underline{\xi}_{100\text{BMR}}(\epsilon)$ is evidenced as $\epsilon$ varies from 0 to 1 in Figure 3. This suggests that the objective $IG(0.001, 0.001)$ prior

is essentially equivalent to an alternative objective $Gamma(0.001, 0.001)$ prior in approximating the improper reciprocal prior, at least for these data. For either prior on $\gamma_0$, the various $\underline{\xi}_{100\text{BMR}}(\epsilon)$ values are also typically smaller than those from the other scenarios: the objective priors appear to consistently shrink the BMDL towards zero.

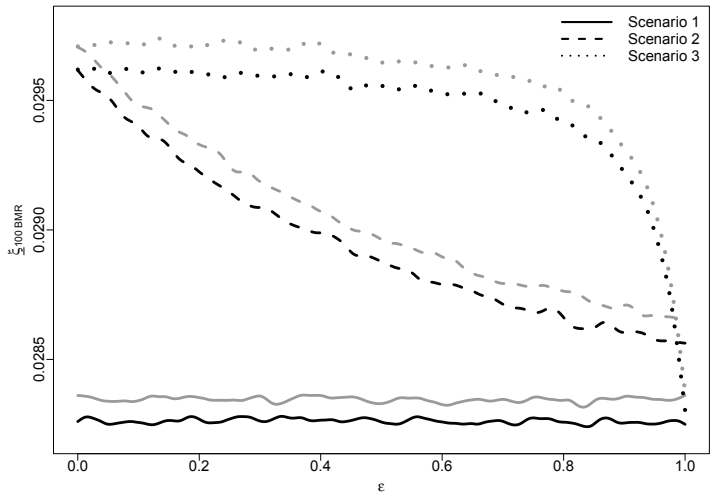


Figure 3: Kernel smoothed value of $\underline{\xi}_{10}$ as a function of changing $\epsilon$ in $\epsilon$-contamination study with cumene carcinogenicity data from Table 1. Solid curves (———) are from Scenario 1: an objective $IG(0.001, 0.001)$ prior for $\xi$ is contaminated by an objective $Gamma(0.001, 0.001)$ prior. Dashed curves (- - - - -) are from Scenario 2: the elicited IG prior for $\xi$ is contaminated by the elicited gamma prior. Dotted curves ($\cdots\cdots$) are from Scenario 3: the elicited IG prior for $\xi$ is contaminated by an objective $Gamma(0.001, 0.001)$ prior. Gray curves indicate an objective $Beta\left(\frac{1}{2}, \frac{1}{2}\right)$ prior for $\gamma_0$; black curves indicate the elicited beta prior for $\gamma_0$. Dose scale ($\xi$) is standardized to unit length after dividing by the highest dose in the data set.

In Scenario 2, $\underline{\xi}_{100\text{BMR}}(\epsilon)$ decreases in an exponentially decaying fashion, to a maximum drop of about $3 - 4\%$ as $\epsilon$ increases from 0 to 1 (Figure 3). This smooth decrease suggests that the contaminating gamma prior affects $\underline{\xi}_{100\text{BMR}}$ in a consistent fashion as $\epsilon$ changes. However, the relatively small maximum decrease also suggests that the sensitivity of $\underline{\xi}_{100\text{BMR}}$ to the contaminating elicited gamma is limited.

In Scenario 3, small decreases in $\underline{\xi}_{100\text{BMR}}(\epsilon)$ are evidenced as $\epsilon$ increases from 0 to approximately 0.8, followed by a precipitous drop as $\epsilon$ increases to 1 (Figure 3). The maximum decrease is slightly more than 4%, however. This provides some evidence for the robustness of $\underline{\xi}_{100\text{BMR}}$ under an elicited IG prior; a significant amount of contamination from the objective prior is required to greatly reduce the BMCL. In each scenario, we also see that these patterns occur for either form of prior for $\gamma_0$, although, on average, the elicited $\gamma_0$ prior consistently produces a roughly 0.3% smaller BMCL than the objective $\gamma_0$ prior.

Table 3 gives the values for $|D(q)|$ under each scenario. We also include a relative error,

measuring the maximum change in $\underline{\xi}_{100\text{BMR}}(\epsilon)$ compared to $\underline{\xi}_{100\text{BMR}}(0)$:

$$\delta_{100\text{BMR}} = \frac{\underline{\xi}_{100\text{BMR}}(0) - \min\left\{\underline{\xi}_{100\text{BMR}}(\epsilon)\right\}}{\underline{\xi}_{100\text{BMR}}(0)}.$$

The relative errors for Scenario 1 are all less than 1%, suggesting little overall decrease in $\underline{\xi}_{100\text{BMR}}$ as $\epsilon$ increased from 0 to 1. For Scenarios 2 and 3, the relative errors are all approximately 4% (also noted above). Although non-trivial, such small changes suggests only minor sensitivity of the base priors to the contaminating priors, at least for these data. Values of $|D(q)|$ for Scenario 1 and 3 are all very close to 0, suggesting tiny instantaneous change of $\underline{\xi}_{100\text{BMR}}$ at $\epsilon = 0$. These are consistent with the graphical patterns seen in Figure 3. By contrast, for Scenario 2 $|D(q)|$ is roughly 1 to 3 orders of magnitudes higher than the other scenarios. This is again consistent with the patterns in Figure 3.

Table 3: Relative errors, $\delta_{100\text{BMR}}$, and instantaneous change measure, $|D(q)|$, for each prior contamination scenario (see text) with the cumene carcinogenicity data in Table 1.

| | Objective beta prior | | Elicited beta prior | |
|---|---|---|---|---|
| Scenario | $\delta_{100\text{BMR}}$ | $\lvert D(q)\rvert$ | $\delta_{100\text{BMR}}$ | $\lvert D(q)\rvert$ |
| 1 | $7.679 \times 10^{-3}$ | $1.134 \times 10^{-4}$ | $5.287 \times 10^{-3}$ | $5.632 \times 10^{-6}$ |
| 2 | $3.767 \times 10^{-2}$ | $1.905 \times 10^{-3}$ | $3.645 \times 10^{-2}$ | $1.839 \times 10^{-3}$ |
| 3 | $4.477 \times 10^{-2}$ | $4.536 \times 10^{-5}$ | $4.396 \times 10^{-2}$ | $4.174 \times 10^{-5}$ |

## 4.3 Posterior Extra Risks

The posterior extra risks are further compared, according to the description in §3.5.1. For a comparison, we fix two crucial dose values: our 95% Bayesian BMCL of $d = \underline{\xi}_{10} = 14.752$ ppm and the frequentist 95% BMCL of $d = \underline{\xi}_{10} = 13.618$ ppm. Figure 4 presents kernel density estimates of the posterior extra risk at each $d$. Both density estimates appear roughly symmetric and lie fairly close to each other. As expected, since it is taken at a higher level of dose, the posterior density at the Bayesian BMCL locates farther up the extra risk scale: the mean extra risk at the Bayesian BMCL is 0.083 while the mean extra risk at the frequentist BMCL is 0.077. This reiterates the greater conservatism of the frequentist estimates with these data.

The standard deviation of the posterior extra risk at the Bayesian BMCL is 0.0096 while the standard deviation of the posterior extra risk at the frequentist BMCL is 0.0090. Again, these values are comparable; the frequentist BMCL is slightly less variable. The 95th percentile of the extra risks at the Bayesian BMCL is exactly 0.1, as expected; the 95th percentile of the extra risks at the frequentist BMCL is 0.093 which again illustrates the conservatism of the frequentist BMCL.

Figure 5 presents a 95% posterior credible band for the extra risk, using the method described in §3.5.1. The centroid estimate for the extra risk is also included (dashed curve). By construction, the corresponding dose level at BMR = 0.1 on the 95% band coincides with the Bayesian BMCL obtained above.

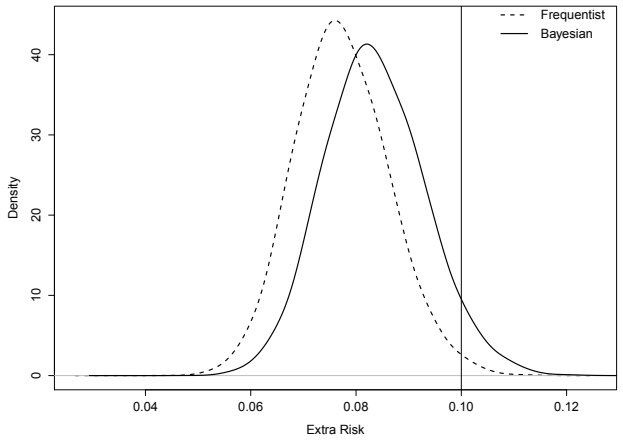


Figure 4: Kernel density estimators of posterior distribution for extra risk $R_E(d)$, where $d$ is taken as the Bayesian BMCL (———) and as the frequentist BMCL (- - - - -) in Example 1. Vertical bar indicates the BMR at 0.10.

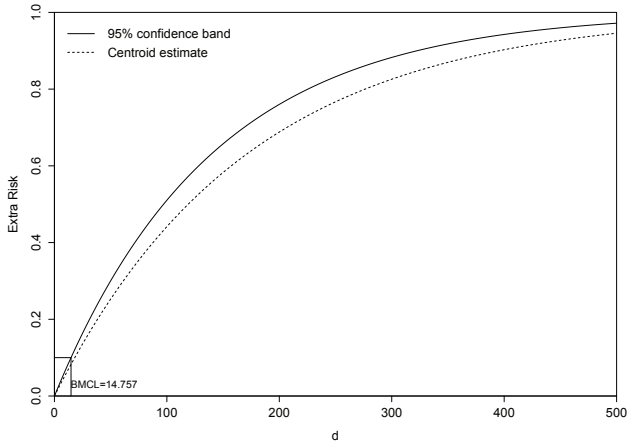


Figure 5: 95% posterior credible band (———) and centroid estimate (- - - - -) for extra risk function $R_E(d)$ in Example 1. Bayesian BMCL at BMR = 0.1 also marked, via inversion of the 95% band.

# 5 Discussion

Herein, we consider a Bayesian approach for estimating benchmark doses (BMDs) in quantitative risk analysis. Placing emphasis on environmental carcinogenicity assessment, our method estimates the BMD via meaningfully reparameterized quantal-response models. Prior information for the parameters is incorporated through an elicitation process, although use of objective priors is also considered. Due to the complexity of the reparameterized models, the joint posterior distribution for the unknown parameters is approximated via Monte Carlo methods, using a computationally intensive but stable adaptive Metropolis (AM) algorithm. A lower Bayesian credible limit (BMDL) is estimated as the lower percentile of the AM sample. Environmental risk analysts can employ this Bayesian approach to construct inferences on the BMD/BMDL by incorporating expert prior knowledge for the model parameters.

Of course, some caveats and qualifications are in order. Objective prior specification can be approached via many strategies when elicitation breaks down. We have assumed independent $IG(0.001, 0.001)$ and $Beta\left(\frac{1}{2}, \frac{1}{2}\right)$ priors for $\xi$ and $\gamma_0$, respectively. Other forms may be pertinent, however, and we acknowledge the possible subjectivity that these choices bring to our model.

Another important issue concerns potential *model uncertainty*. Our representation for the risk function was based on the popular quantal linear model. Many other forms, some quite complex, can be chosen to model $R(d)$, however. The logistic form $R(d) = 1/(1 + e^{-\beta_0 - \beta_1 d})$ from §1.2 is a highly popular alternative (Foronda *et al.*, 2007; Shao and Small, 2012). Or, the quantal linear model can be extended into a 'two-stage' model, $R(d) = 1 - \exp(-\beta_0 - \beta_1 d - \beta_2 d^2)$, where $\beta_j \geq 0 \; \forall j = 0, 1, 2$. Indeed, it is straightforward to employ our reparameterization strategy with both these models: all aspects of the Bayesian estimation scheme in §3 remain valid, including our elicitation strategies for the priors and the AM approach for building the joint posterior. [The three $\beta_j$s in the two-stage model require incorporation of an additional, reparameterized, nuisance parameter. Details are given in Fang (2014).]

For instance, if we apply the logistic model to the cumene carcinogenicity data in Table 1, the BMC estimates on the original ppm scale are $\hat{\xi}_{10}$=42.946 ppm (if using the sample posterior median) or $\hat{\xi}_{10}$=40.892 ppm (if using the sample posterior lower tercile). The 95% BMCL is $\underline{\xi}_{10}$=35.599 ppm. These values contrast greatly with those from the quantal linear model in §4.1, leading one to ask which is more appropriate—a classic example of model uncertainty. To compare the two fits, we turn to the Bayes Factor ($BF$) (Kass and Raftery, 1995): the $BF$ comparing the quantal linear model to the logistic model is calculated as the ratio of corresponding marginal likelihoods:

$$BF = \frac{m_{\text{ql}}(\boldsymbol{Y})}{m_{\text{lo}}(\boldsymbol{Y})}, \tag{5.1}$$

where $m_{\text{ql}}(\boldsymbol{Y})$ denotes the marginal likelihood using the quantal linear model and $m_{\text{lo}}(\boldsymbol{Y})$ denotes the marginal likelihood using the logistic model. We approximate both marginal likelihoods using the geometric estimator and bridge sampling method from §3.4, producing $BF = 518.3$ for the cumene data. Following Kass and Raftery (1995, §3.2), since this $BF$ exceeds 150 we say there is 'very strong' evidence that the quantal linear model fits the data better than the logistic model. (This is, perhaps, not surprising: the concave response observed in Figure 2 would be difficult for a logistic model to match, compared to the always-concave quantal linear form.)

Clearly, model adequacy is an important issue in benchmark risk analysis (West *et al.*, 2012). It has received limited attention, however, particularly from a Bayesian perspective (Shao and Small, 2012). We are expanding our Bayesian BMD estimation approach to consider other popular quantal-response models, and to deal with model adequacy concerns. We hope to report on this in a future manuscript.

# Acknowledgements

Thanks are due Drs. Anton Westveld, D. Dean Billheimer, Susan J. Simmons, and Cuixian (Tracy) Chen for their helpful suggestions during the preparation of this material. The results represent a portion of the first author's Ph.D. dissertation with the University of Arizona Graduate Interdisciplinary Program in Statistics.


# Appendix

In this appendix we summarize the empirical screen employed to mark quantal data sets as dose-response 'data failures.' Over increasing doses $0 = d_1 < \ldots < d_i < \ldots < d_m$ $(i = 1, \ldots m)$, we first calculate the empirical extra risks:

$$\tilde{R}_E(d_i) = \frac{\frac{Y_i}{N_i} - \frac{Y_1}{N_1}}{1 - \frac{Y_1}{N_1}}.$$

We then connect each point $\left(d_i, \tilde{R}_E(d_i)\right)$ to the origin and find the largest slope, $S_{\max}$, among all these $m-1$ rays. If $S_{\max} \leq 0$, no increasing trend is evidenced and we mark this as a 'data failure'. Notice that we do not perform a formal trend test to detect a decreasing trend (Wheeler and Bailer, 2009a), but we do require that at least one estimated risk for some $d_i$ $(i > 1)$ is higher than the estimated background risk.

# References


Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18**, 343–383.

Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.

Buckley, B. E., Piegorsch, W. W. and West, R. W. (2009). Confidence limits on one-stage model parameters in benchmark risk assessment. *Environmental and Ecological Statistics* **16**, 53–62.

Casella, G. and Berger, R. L. (2002). *Statistical Inference*, 2nd edn. Pacific Grove, CA: Duxbury.

Christensen, R., Johnson, W. O., Branscum, A. J. and Hanson, T. E. (2011). *Bayesian Ideas and Data Analysis: An Introduction for Scientists and Statisticians*. Boca Raton, FL: Chapman & Hall/CRC Press.

Coull, B. A., Mezzetti, M. and Ryan, L. M. (2003). A Bayesian hierarchical model for risk assessment of methylmercury. *Journal of Agricultural, Biological, and Environmental Statistics* **8**, 253–270.

Crump, K. S. (1984). A new method for determining allowable daily intake. *Fundamental and Applied Toxicology* **4**, 854–871.

Crump, K. S. (1995). Calculation of benchmark doses from continuous data. *Risk Analysis* **15**, 79–89.

Davis, J. A., Gift, J. S. and Zhao, Q. J. (2012). Introduction to benchmark dose methods and U.S. EPA's Benchmark Dose Software (BMDS) version 2.1.1. *Toxicology and Applied Pharmacology* **254**, 181–191.

European Union (2003). *Technical Guidance Document (TGD) on Risk Assessment of Chemical Substances following European Regulations and Directives, Parts I-IV*. Technical Report number EUR 20418 EN/1-4. Ispra, Italy: European Chemicals Bureau (ECB).

Fang, Q. (2014). *Hierarchical Bayesian Benchmark Risk Analysis*. Ph.D. thesis, Interdisciplinary Program in Statistics, University of Arizona, Tucson, AZ.

Foronda, N. M., Fowles, J., Smith, N., Taylor, M. and Temple, W. (2007). A benchmark dose analysis for sodium monofluoroacetate (1080) using dichotomous toxicity data. *Regulatory Toxicology and Pharmacology* **47**, 84–89.

Gelman, A. (2012). Posterior distribution. In El-Sharaawi, A. H. and Piegorsch, W. W. (eds.), *Encyclopedia of Environmetrics*, 2nd edn., **4**, 2020–2021. Chichester: John Wiley & Sons.

Guha, N., Roy, A., Kopylev, L., Fox, J., Spassova, N., and White, P. (2013). Nonparametric Bayesian methods for benchmark dose estimation. *Risk Analysis* **9**, 1608–1619.

Held, L. (2004). Simultaneous inference in risk assessment; a Bayesian perspective. In Antoch, J. (ed.), *COMPSTAT 2004. Proceedings in Computational Statistics*, 213–222. Heidelberg: Physica-Verlag.

Kass, R. E. and Raftery, A. E. (1995). Bayes Factor. *Journal of the American Statistical Association* **90**, 773–795.

Kuhnert, P. M. (2011). Four case studies in using expert opinion to inform priors. *Environmetrics* **22**, 662–674.

Kvam, P. H. and Vidakovic, B. (2007). *Nonparametric Statistics with Applications to Science and Engineering*. Hoboken, NJ: John Wiley & Sons.

Lambert, P. C., Sutton, A. J., Burton, P. R., Abrams, K. R. and Jones, D. R. (2005). How vague is vague? A simulation study of the impact of the use of vague prior distributions in MCMC using WinBUGS. *Statistics in Medicine* **24**, 2401–2428.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica* **14**, 41–67.

Meng, X. and Wong W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica* **6**, 831–860.

Morales, K. H., Ibrahim, J. G., Chen, C.-J. and Ryan, L. M. (2006). Bayesian model averaging with applications to benchmark dose estimation for arsenic in drinking water. *Journal of the American Statistical Association* **101**, 9–17.

Naufal, Z., Kathman, S. and Wilson, C. (2009). Bayesian derivation of an oral cancer slope factor distribution for 4-(methylnitrosamino)-1-(3-pyridyl)-1-butanone (NNK). *Regulatory Toxicology and Pharmacology* **55**, 69–75.

Nitcheva, D. K., Piegorsch, W. W. and West, R. W. (2007). On use of the multistage dose-response model for assessing laboratory animal carcinogenicity. *Regulatory Toxicology and Pharmacology* **48**, 135–147.

OECD (2006). *Current Approaches in the Statistical Analysis of Ecotoxicity Data: A Guidance to Application, Series on Testing and Assessment No. 54.* Paris: Environment Directorate, Organisation For Economic Co-Operation and Development.

OECD (2008). *Draft Guidance Document on the Performance of Chronic Toxicity and Carcinogenicity Studies, Supporting TG 451, 452 and 453.* Paris: Organisation For Economic Co-Operation and Development.

O'Hagan, A. (1994). *Kendall's Advanced Theory of Statistics, Volume 2B, Bayesian Inference*, 2nd edn. London: Edward Arnold.

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E. and Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts' Probabilities.* Chichester: John Wiley & Sons.

Parham, F. and Portier C. J. (2005). Benchmark dose approach. In Edler, L. and Kitsos, C. (eds.), *Recent Advances in Quantitative Methods in Cancer and Human Health Risk Assessment*, 239–254. Chichester: John Wiley & Sons.

Piegorsch, W. W. and Bailer, A. J. (2005). *Analyzing Environmental Data.* Chichester: John Wiley & Sons.

R Development Core Team. (2012). *R: A Language and Environment for Statistical Computing.* Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Shao, K. (2012). A comparison of three methods for integrating historical information for Bayesian model averaged benchmark dose estimation. *Environmental Toxicology and Pharmacology* **34**, 288–296.

Shao, K. and Small, M. J. (2011). Potential uncertainty reduction in model-averaged benchmark dose estimates informed by an additional dose study. *Risk Analysis* **31**, 1561–1575.

Shao, K. and Small, M. J. (2012). Statistical evaluation of toxicological experimental design for Bayesian model averaged benchmark dose estimation with dichotomous data. *Human and Ecological Risk Assessment* **18**, 1096–1119.

Stern, A. H. (2008). Environmental health risk assessment. In Melnick, E. L. and Everitt, B. S. (eds.), *Encyclopedia of Quantitative Risk Analysis and Assessment* **2**, 580–589. Chichester: John Wiley & Sons.

U.S. EPA (2012). *Benchmark Dose Technical Guidance Document.* Technical Report number EPA/100/R-12/001. Washington, DC: U.S. Environmental Protection Agency.

U.S. General Accounting Office (2001). *Chemical Risk Assessment. Selected Federal Agencies' Procedures, Assumptions, and Policies.* Report to Congressional Requesters number GAO-01-810. Washington, DC: U.S. General Accounting Office.

U.S. National Toxicology Program (2009). *Toxicology and Carcinogenesis Studies of Cumene (CAS NO. 98-82-8) in F344/N Rats and B6C3F$_1$ Mice.* Technical Report number 542. Research Triangle Park, NC: U.S. Department of Health and Human Services, Public Health Service.

West, R. W., Piegorsch, W. W., Peña, E. A., An, L., Wu, W., Wickens, A. A., Xiong, H., and Chen, W. (2012). The impact of model uncertainty on benchmark dose estimation. *Environmetrics* **23**, 706-716.

Wheeler, M. W. and Bailer, A. J. (2009a). Comparing model averaging with other model selection strategies for benchmark dose estimation. *Environmental and Ecological Statistics* **16**, 37–51.

Wheeler, M. W. and Bailer, A. J. (2009b). Benchmark dose estimation incorporating multiple data sources. *Risk Analysis* **29**, 249–256.

Wheeler, M. W. and Bailer, A. J. (2012). Monotonic Bayesian semiparametric benchmark dose analysis. *Risk Analysis* **32**, 1207–1218.

arXiv:1406.2725v1 [stat.ME] 10 Jun 2014

# Supplement to 'Bayesian Benchmark Dose Analysis'

**Qijun Fang**[a†], **Walter W. Piegorsch**[a,b], and **Katherine Y. Barnes**[a,c]
[a]*Program in Statistics*, [b]*BIO5 Institute*, and [c]*James E. Rogers College of Law*
*University of Arizona, Tucson, AZ, USA*

This supplementary document provides supporting material for *Bayesian Benchmark Dose Analysis* by Qijun Fang and Walter W. Piegorsch (the 'main document'). The various sections below address a variety of supplemental/supporting topics, and are not intended to flow naturally between each other. They are, however, presented in roughly the same order in which their counterpart topics appear in the main document. As there, we denote the Benchmark Dose (BMD) target parameter as $\xi$ and the background response probability nuisance parameter as $R(0) = \gamma_0$.

## 1 Finding prior parameters from elicited quantiles

We give here technical aspects on derivation of the prior parameters $\alpha$, $\beta$, $\psi$, and $\omega$ for the prior densities $\xi \sim IG(\alpha, \beta)$ and $\gamma_0 \sim Beta(\psi, \omega)$ used in the main document's hierarchical model. The elicited lower quartiles and medians $Q_{1\xi}, Q_{2\xi}$ for $\xi$, and $Q_{1\gamma}, Q_{2\gamma}$ for $\gamma_0$, respectively, are assumed given from domain expert(s) input.

Start with the elicitation for $\xi$: by definition, $Q_{1\xi}$ and $Q_{2\xi}$ satisfy

$$\int_0^{Q_{j\xi}} \frac{\beta^\alpha}{\Gamma(\alpha)} \xi^{(-\alpha-1)} \exp\left\{-\frac{\beta}{\xi}\right\} d\xi = \frac{j}{4},$$

$(j = 1, 2)$, where the integrand is the IG p.d.f. This establishes a system of non-linear equations for $\alpha$ and $\beta$. Unfortunately, no closed-form solution exists for the system, and so we turn to numerical methods. We employ a gradient method discussed by Barzilai and Borwein (1988) and implemented in the `R` statistical environment (R Development Core Team, 2012) via the `BB` package (Varadhan and Gilbert, 2009). This method employs iterative updating until half the $L_2$ norm of the system at the proposed solution is smaller than $10^{-10}$.

Similarly, for $\gamma_0$ the elicited quartiles $Q_{1\gamma}$ and $Q_{2\gamma}$ satisfy

$$\int_0^{Q_{j\gamma}} \frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)} \gamma_0^{\psi-1}(1-\gamma_0)^{\omega-1} d\gamma_0 = \frac{j}{4},$$

$(j = 1, 2)$, now applying the p.d.f. for the beta prior. Here again, the resulting system of non-linear equations possesses no closed-form solution, so we appeal to the `BB` package in `R`.

The resulting iterative solutions to these two systems of equations produce values for the prior parameters $\alpha$, $\beta$, $\psi$, and $\omega$ which are then employed in our hierarchical model for the BMD.

[†]Email address: qijunf@email.arizona.edu; Corresponding author

## 2 The adaptive Metropolis algorithm

The posterior p.d.f. under our hierarchical model is

$$\begin{aligned}\pi(\xi,\gamma_0|\boldsymbol{Y}) &= \frac{f(\boldsymbol{Y}|\xi,\gamma_0)\pi(\xi|\alpha,\beta)\pi(\gamma_0|\psi,\omega)}{m(\boldsymbol{Y})} \\ &= \frac{\prod_{i=1}^{n} R(d_i)^{Y_i}[1-R(d_i)]^{N_i-Y_i}}{m(\boldsymbol{Y})}\frac{\beta^\alpha e^{-\beta/\xi}}{\Gamma(\alpha)\xi^{\alpha+1}}\frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)}\gamma_0^{\psi-1}(1-\gamma_0)^{\omega-1},\end{aligned} \tag{2.1}$$

where

$$m(\boldsymbol{Y}) = \int_0^\infty \int_0^1 f(\boldsymbol{Y}|\xi,\gamma_0)\pi(\xi|\alpha,\beta)\pi(\gamma_0|\psi,\omega)d\gamma_0 d\xi.$$

By employing proper priors, this is guaranteed to be integrable (Gelman *et al.*, 2004, §5.3), which helps motivate our preference for the IG and beta prior assumptions on $\xi$ and $\gamma_0$, respectively. Nonetheless, the posterior is intractable, so we turn to approximation by computer. Markov chain Monte Carlo (McMC) approaches have been applied with great success to models such as this (Robert and Casella, 2011), although the complexity of (2.1) prevents us from applying well-known McMC algorithms such as Gibbs sampling (Gelfand, 2000). Instead, we move to Metropolis-Hastings approaches (Robert and Casella, 2004, Ch. 7) and in particular consider *adaptive* Metropolis (AM) algorithms. These tune the variance of the underlying Metropolis proposal density adaptively when generating ongoing draws of the chain. From our experience with a variety of such methods, we favor a global AM procedure with componentwise adaptive scaling described by Andrieu and Thoms (2008). Their 'AM6' algorithm uses all previous iterations in the chain to update the variance-covariance matrix of the current bivariate proposal density. The operation is global in the sense that both the $\xi$ and $\gamma_0$ parameters are updated simultaneously at each $k$th bivariate draw; $k = 1, \ldots, K$. It is also componentwise in the sense that the updating step sizes for $\xi$ and $\gamma_0$ are controlled by separate adaptive scaling parameters. See Andrieu and Thoms (2008) for further details.

To begin the AM chain, 'starting' points for $\xi$ and $\gamma_0$ must be chosen. Rather than applying a random initialization, we take the starting point for $\gamma_0$ as $\frac{Y_1+0.25}{N_1+0.5}$, i.e., the estimated nonresponse probability with a slight shrinkage term included to pull away from the zero boundary.

The starting point for $\xi$ was obtained by first treating $R_E(d) = S_{\max}d$ as an empirical approximation to the extra risk function and then solving for $d$ in the equation BMR $= S_{\max}d$. Here, $S_{\max}$ is the maximum empirical origin-to-response slope described in the main document's Appendix. We then find the starting point for $\xi$ as simply BMR$/S_{\max}$.

The AM6 algorithm employs a bivariate normal proposal distribution. Andrieu and Thoms (2008) recommend $\frac{2.38^2}{2}\mathbf{I}$ as an initial variance-covariance matrix, where $\mathbf{I}$ is the $2\times 2$ identity matrix. At each subsequent iteration, the variance-covariance matrix of the normal proposal density is then updated based on the previous iteration's variance-covariance matrix and on the current iteration's adaptive scaling parameters for $\xi$ and $\gamma_0$. See Andrieu and Thoms (2008) for details.

# 3 Monte Carlo 'burn-in' diagnostics

To account for potential instability in the early portions of the bivariate chain, we include a 'burn-in' for the Monte Carlo sample (Gelman *et al.*, 2004, §11.6). We couple this with the larger question of how to assess the chain's convergence. Many approaches exist for diagnosing McMC convergence (Cowles and Carlin, 1996), and indeed, the issue is a topic of ongoing debate (Robert and Casella, 2004, §12.2). We mimic a method due to Geweke (1992), where early portions of the chain are sub-sampled and compared against latter portions to determine where the larger sample of $K$ draws begins to exhibit stable performance. To approximate independence between the two sub-samples, we bifurcate the chain by a gap of no less than $\frac{K}{5}$ consecutive draws.

For summary diagnostics, we calculate the difference in arithmetic means between the two bifurcated sub-samples, and divide each difference by its standard error to produce a $Z$-statistic. The approximate standard error is taken as the square root of the sum of the variances of each mean. This is done separately for both the $\xi_k$ and $\gamma_{0k}$ components, $k = 1, \ldots, K$. (For sake of simplicity, we make no correction for multiplicity.) To adjust for possible autocorrelation within each sub-sample, the individual variances are based on estimated spectral densities at frequency zero; see Fang (2013) for full details.

To monitor if the pattern of association between $\xi$ and $\gamma_0$ within the larger chain also exhibits reasonable stability, we include comparison of the covariances between the two sub-samples. As a diagnostic measure here we use the sample covariance from each sub-sample,

$$\widehat{\mathrm{Cov}}(\xi, \gamma_0) = \frac{1}{L} \sum_{k=1}^{L} (\xi_k - \bar{\xi})(\gamma_{0k} - \bar{\gamma}_0),$$

where $\bar{\xi}$ and $\bar{\gamma}_0$ are the pertinent arithmetic means within a bifurcated sub-sample of length $L$. To find the approximate variance of each estimated covariance we calculate the quantities $\psi_k = (\xi_k - \bar{\xi})(\gamma_{0k} - \bar{\gamma_0})$ across all values of $k$ in the given sub-sample and then estimate the variance of $\psi_k$ based on their estimated spectral densities at frequency zero.

Each such comparison is conducted over three individual, separated bifurcations of the full $K$-length sample: (i) the first 10% of the chain ($L = K/10$) vs. the final 50% ($L = K/2$); (ii) the first 20% of the chain ($L = K/5$) vs. the final 50%; and (iii) the first 30% of the chain ($L = 0.3K$) vs. the final 50%. We consider an individual diagnostic comparison as a 'pass' if the corresponding $Z$-statistic comparing the two bifurcated samples is less than 1.96 in absolute value. To 'pass' the full diagnostic at each bifurcation, all three measures—mean of $\xi$, mean of $\gamma_0$, and covariance of $\{\xi,\gamma_0\}$—must individually pass.

The comparisons are performed in sequential order: if the 10%-vs.-50% diagnostic fails, then the 20%-vs.-50% diagnostic is conducted. If this fails, we move to the 30%-vs.-50% diagnostic. When a diagnostic passes, we view the indicated early portion of the chain as a reasonable burn-in. For notation, we let $K_0 < K$ be the resulting index that begins the retained portion of the chain. So, e.g., if the 10%-vs.-50% diagnostic fails but then the 20%-vs.-50% diagnostic passes, we take the first 20% of the chain as burn-in and use the remaining 80% of the chain as our Monte Carlo sample from $\pi(\xi, \gamma_0|\boldsymbol{Y})$. If $K = 100,000$, as used throughout, this gives $K_0 = 20,001$.

If none of the three sequential diagnostics passes, we re-set the random-number generator's seed and reassess a new set of $K$ Monte Carlo draws. If after five such re-starts the diagnostic continues to fail, we report an *algorithm failure*. (This did not occur very often; see the simulation results in §6, below.)

# 4 Estimating marginals via bridge sampling

In the main document, we often have reason to calculate (an approximation to) a marginal likelihood of the form

$$m(\boldsymbol{Y}) = \int_0^\infty \int_0^1 f(\boldsymbol{Y}|\xi,\gamma_0)\pi(\xi|\alpha,\beta)\pi(\gamma_0|\psi,\omega)d\gamma_0 d\xi$$

for various choices of $\pi(\xi|\alpha,\beta)$ and/or $\pi(\gamma_0|\psi,\omega)$. To do so, we construct a *geometric estimator* from our AM sample, appealing to the 'bridge sampling' method recommended by Meng and Wong (1996) and Lopes and West (2004). To approximate the marginal likelihood, an approximation $g(\cdot)$ to the joint posterior density of $(\xi,\gamma_0)$ is first chosen. For simplicity, we took $g(\xi,\gamma_0)$ as the bivariate normal density function with mean set to the empirical mean vector and variance set to the empirical variance-covariance matrix of $\xi$ and $\gamma_0$ estimated from the AM sample. Next, a set of $K^* = K - K_0 + 1$ bivariate vectors $\left\{[\xi_j^* \ \gamma_{0j}^*]^T\right\}_{j=1}^{K^*}$ are drawn from $g(\xi,\gamma_0)$. The marginal likelihood is then approximated by substituting the (retained) AM sample, $\left\{[\xi_k,\gamma_{0k}]^T\right\}_{k=K_0}^{K}$, and the generated sample from $g(\xi,\gamma_0)$, $\left\{[\xi_j^* \ \gamma_{0j}^*]^T\right\}_{j=1}^{K^*}$, into the ratio

$$\widehat{m}(\boldsymbol{Y}) = \frac{\frac{1}{K^*}\sum_{j=1}^{K^*}\left\{\pi\left(\xi_j^*|\alpha,\beta\right)\pi\left(\gamma_{0j}^*|\psi,\omega\right)f\left(\boldsymbol{Y}|\xi_j^*,\gamma_{0j}^*\right)/g\left(\xi_j^*,\gamma_{0j}^*\right)\right\}^{1/2}}{\frac{1}{K^*}\sum_{k=K_0}^{K}\left[g\left(\xi_k,\gamma_{0k}\right)/\left\{\pi\left(\xi_k|\alpha,\beta\right)\pi\left(\gamma_{0k}|\psi,\omega\right)f\left(\boldsymbol{Y}|\xi_k,\gamma_{0k}\right)\right\}\right]^{1/2}},$$

where $f(\cdot)$ denotes the binomial likelihood function and $\pi(\cdot)$ denotes the pertinent prior density [e.g., $q(\xi)$ when approximating $m_q(\boldsymbol{Y})$ in our $\epsilon$-contamination study].

# 5 Prior elicitation with the cumene data

To elicit the prior parameters for the cumene data in Example 1, we applied input from domain experts and based the prior elicitation on existing background from the toxicological literature. We began with the BMD target parameter, $\xi$: an oral No Observed Adverse Effect Level (NOAEL) for long-term cumene exposure in rodents was given by the EPA IRIS website (`http://www.epa.gov/iris/subst/0306.htm`) as 110 mg/kg-day. Guidance for converting the oral mg/kg-day metric to inhalation ppm (the scale used in the original data table from the main document) was available via conversion data given by the California OEHHA (2004): for rodents, 1 mg/kg-day works out to roughly 2.25 ppm. Arguing that a NOAEL is loosely equivalent to a BMD—although, see (Kodell, 2005) for an in-depth discussion—the median prior estimate was $Q_{2\xi}$ = 110 mg/kg-day = 247.5 ppm $\approx$ 250 ppm. For the first quartile, the EPA IRIS site gave a cumene inhalation NOAEL for rodents as 435 mg/cu.m,

based on a shorter, 13-week study. The shorter exposure period was suited to provide only a limited estimate on the BMD's central tendency; we translated this to use the 435 mg/cu.m 13-week NOAEL as a prior estimate of the first quantile. For conversion to inhalation ppm, the EPA indicated that roughly 1 mg/cu.m = 0.204 ppm, so we took $Q_{1\xi} = 435$ mg/cu.m = 88.75 ppm $\approx 90$ ppm. We then divided the ppm dose by 500 ppm to standardize the $\xi$-scale to $0 \leq \xi \leq 1$. This produced prior first quartile and median estimates for $\pi(\xi|\alpha, \beta)$ as 0.18 and 0.5, respectively.

For $\gamma_0$, historical control data on alveolar/bronchiolar tumors in B6C3F$_1$ mice were given by U.S. NTP (2009, Table D3a). These provide direct prior information on the background risk. The historical data gave a median background response of $Q_{2\gamma} = 0.08$ and a lower quartile of $Q_{1\gamma} = 0.04$.

These quartile-based elicitations were then used to construct the prior densities $\xi \sim IG(\alpha, \beta)$ and $\gamma_0 \sim Beta(\psi, \omega)$ as described in §1, above.

# 6 Trace plot and convergence diagnostics for cumene example

In the cumene example, an AM approach was applied to approximate the posterior for $\xi$ and $\gamma_0$. At an initial chain size of $100,000$, Figure 1 displays the full trace plot for the $\xi$ component of this AM sample. The plot exhibits stable performance and fairly fast mixing. Our convergence diagnostic procedure recommended an initial burn-in of $10,000$ draws. The remaining $90,000$ draws were then used as the posterior approximation for these data.

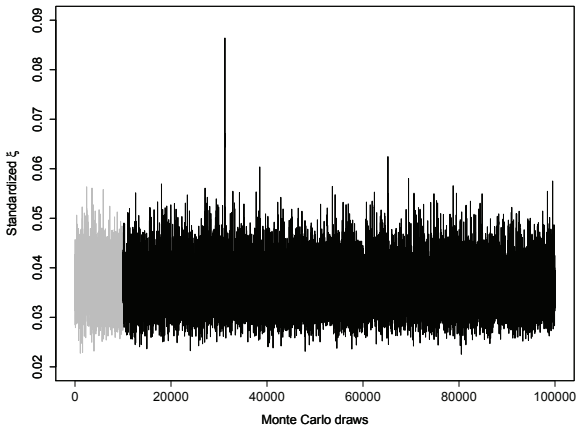


Figure 1: Trace plot for $\xi$ in Example 1. Dose scale ($\xi$) is standardized to unit length after dividing by the highest dose in the data set. First 10,000 Monte Carlo draws in grey indicate burn-in period.

# References


Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18**, 343–383.

Barzilai, J. and Borwein, J. M. (1988). Two-point step size gradient methods. *IMA Journal of Numerical Analysis* **8**, 141–148.

California Office of Environmental Health Hazard Assessment (2004). *Proposition 65 Maximum Allowable Dose Level (MADL) for Reproductive Toxicity for Methyl Bromide as a Structural Fumigant*. Sacramento, CA: CA Office of Environmental Health Hazard Assessment (OEHHA), Reproductive and Cancer Hazard Assessment Section.

Cowles, M. K. and Carlin, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: A comprarative review. *Journal of the American Statistical Association*, **91**, 883–904.

Fang, Q. (2013). *Hierarchical Bayesian Benchmark Risk Analysis*. Ph.D. thesis, Interdisciplinary Program in Statistics, University of Arizona, Tucson, AZ.

Gelfand, A. E. (2000). Gibbs sampling. *Journal of the American Statistical Association* **95**, 1300–1304.

Gelman, A., Carlin, B. J., Stein, H. S. and Rubin, D. B. (2004). *Bayesian Data Analysis*, 2nd edn. Boca Raton, FL: Chapman & Hall/CRC.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In Bernardo, J. M., Berger, J. O., Dawid, A. P. and Smith, A. F. M. (eds.), *Bayesian Statistics* **4**, 169–193. Oxford: Oxford University Press.

Kodell, R. L. (2005). Managing uncertainty in health risk assessment. *International Journal of Risk Assessment and Management* **14**, 193–205.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica* **14**, 41–67.

Meng, X. and Wong W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica* **6**, 831–860.

Piegorsch, W. W., An, L., Wickens, A. A., West, R. W. Peña, E. A. and Wu, W. (2013). Information-theoretic model-averaged benchmark dose analysis in environmental risk assessment. *Environmetrics*, **24**, 143–157.

R Development Core Team. (2012). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*, 2nd Edn. New York: Springer-Verlag.

Robert, C. P. and Casella, G. (2011). A history of Markov chain Monte Carlo: subjective recollections from incomplete data. *Statistical Science* **26**, 102–115.

U.S. National Toxicology Program (2009). *Toxicology and Carcinogenesis Studies of Cumene (CAS NO. 98-82-8) in F344/N Rats and B6C3F$_1$ Mice*. Technical Report number 542. Research Triangle Park, NC: U.S. Department of Health and Human Services, Public Health Service.

Varadhan, R. and Gilbert, P. (2009). BB: An R package for solving a large system of nonlinear equations and for optimizing a high-dimensional nonlinear objective function. *Journal of Statistical Software* **32**, 1–26.